

Express Mailing No.: ER211527875US

PATENT APPLICATION
IBM Docket No.: SJO9-2003-0067US1
Kunzler & Associates Docket No.: 1200.2.96

UNITED STATES PATENT APPLICATION

of

JOSEPH J. FATULA, JR.

for

**APPARATUS, SYSTEM, AND METHOD FOR
AUTONOMIC CONTROL OF GRID SYSTEM RESOURCES**

APPARATUS, SYSTEM, AND METHOD FOR AUTONOMIC CONTROL OF GRID SYSTEM RESOURCES

BACKGROUND OF THE INVENTION

FIELD OF THE INVENTION

[0001] This invention relates to grid computing and more particularly relates to autonomic control of performance resources in a grid computing system.

DESCRIPTION OF THE RELATED ART

[0002] Grid computing is a relatively new technology in the computing industry. Many current grid computing systems are designed to allow multiple interconnected computers, or nodes, to work together on large grid applications, such as computational problems, that would be impossible to do on a single machine. In order to feasibly manage such overwhelming amounts of computational processing, the computational problem may be divided into several smaller, more manageable jobs. This type of shared processing of grid applications is possible in part because of the enforced or incidental idle processor time of many personal and business computers.

[0003] Additionally, grid computing systems may allow for shared storage across a plurality of nodes within the system. In this way, data may be stored remotely on the system, in the form of a backup, for example, and allow large amounts of data to be virtually stored by a single node.

[0004] One of the major issues in designing and implementing grid computing systems is resource management. Managing the performance resources available to the grid computing system is fundamental to the efficacy of the system. For example, with various disparate storage, memory, and processing resources available from the many nodes in a grid computing system, somehow these performance resources must be utilized within the grid system in an organized and productive manner.

[0005] Conventional grid systems typically implement a stimulus-response environment in which operations and commands are executed in simply in response to a given stimulus. For example, if a grid application requires certain processor capacities or a particular amount of storage space, then that amount of processor capacity or storage space may simply be assigned for execution of the grid application. This type of stimulus-response environment does not account for the dynamic nature of distributed computing, in which many performances resources may be dynamically allocated, reclaimed, reserved, and so forth. Due to an inability to regulate grid application operations within such a dynamic environment, grid computing operations must be kept relatively simple to conform to a simple stimulus-response environment.

[0006] Consequently, a need exists for an apparatus, system, and method that facilitate favorable autonomic control of performance resources on a grid computing system. Beneficially, such an apparatus, system, and method would overcome the current limitations on grid applications and grid computing systems.

BRIEF SUMMARY OF THE INVENTION

[0007] The present invention has been developed in response to the present state of the art, and in particular, in response to the problems and needs in the art that have not yet been fully solved by currently available grid computing systems. Accordingly, the present invention has been developed to provide an apparatus, system, and method for autonomic management of grid system resources on a grid computing system that overcome many or all of the above-discussed shortcomings in the art.

[0008] The apparatus for autonomic management of grid system resources on a grid computing system is provided with a logic unit containing a plurality of modules configured to functionally execute the necessary steps of for autonomic management of grid system resources on a grid computing system. These modules in the described embodiments include a global monitor module, a global policy module, a global regulation module, a global reservation module, a global termination module, a global arbitration module, and a global profile module.

[0009] A system of the present invention is also presented for autonomic management of grid system resources on a grid computing system. The system may be embodied in a local area network, a wide area network, a combination of local and wide area networks, one or more wireless networks, an internet-based grid computing network, or any other number of grid computing environments. In particular, the system, in one embodiment, includes a global autonomic management apparatus, a local autonomic management apparatus, a policy module, and a regulation module. The system may further include a subscription manager configured to manage a fee subscription for each of the clients connected to the grid computing system.

[0010] A client is also presented for autonomic control of a grid system resource on a grid computing system. In one embodiment, the client is provided with a logic unit containing a plurality of modules configured to functionally execute the necessary steps of autonomic control of a grid system resource on a grid computing system. These modules in

the described embodiments include a client monitor module, a client policy module, a client notification module, a client allocation module, a client reclamation module, a client initiation module, a client termination module, and a client profile module.

[0011] A method of the present invention is also presented for autonomic management of grid system resources on a grid computing system. The method in the disclosed embodiments substantially includes the steps necessary to carry out the functions presented above with respect to the operation of the described apparatus and system. In one embodiment, the method includes monitoring the grid computing system for a trigger event, accessing one of a plurality of system policies, regulating the system resource in response to a recognized trigger event according to one of the plurality of system policies. In a certain embodiment, each of the plurality of system policies corresponding to an operational control parameter of a system resource of the grid computing system. Furthermore, the trigger event may be an initiation trigger event, a regulation trigger event, or a prediction trigger event.

[0012] The method also may include reserving the system resource for a grid system operation, terminating a reservation of a system resource for a grid system operation, arbitrating conflicting grid system operations according to an arbitration policy, and storing a system resource profile. The system resource profile, in one embodiment, identifies a client resource that is allocated to the grid system.

[0013] One embodiment of the present invention beneficially allows dynamically allocated performance resources to be autonomically managed within a grid computing environment. Additionally, another embodiment of the present invention beneficially allows clients to autonomically allocate and reclaim performance resources within the grid computing system. A further embodiment beneficially enables the global autonomic management apparatus to autonomically reserve allocated performance resources for specific grid applications and operations.

[0014] Reference throughout this specification to features, advantages, or similar language does not imply that all of the features and advantages that may be realized with the

present invention should be or are in any single embodiment of the invention. Rather, language referring to the features and advantages is understood to mean that a specific feature, advantage, or characteristic described in connection with an embodiment is included in at least one embodiment of the present invention. Thus, discussion of the features and advantages, and similar language, throughout this specification may, but do not necessarily, refer to the same embodiment.

[0015] Furthermore, the described features, advantages, and characteristics of the invention may be combined in any suitable manner in one or more embodiments. One skilled in the relevant art will recognize that the invention can be practiced without one or more of the specific features or advantages of a particular embodiment. In other instances, additional features and advantages may be recognized in certain embodiments that may not be present in all embodiments of the invention.

[0016] These features and advantages of the present invention will become more fully apparent from the following description and appended claims, or may be learned by the practice of the invention as set forth hereinafter.

BRIEF DESCRIPTION OF THE DRAWINGS

[0017] In order that the advantages of the invention will be readily understood, a more particular description of the invention briefly described above will be rendered by reference to specific embodiments that are illustrated in the appended drawings. Understanding that these drawings depict only typical embodiments of the invention and are not therefore to be considered to be limiting of its scope, the invention will be described and explained with additional specificity and detail through the use of the accompanying drawings, in which:

[0018] Figure 1 is a schematic block diagram illustrating one embodiment of a grid system in accordance with the present invention;

[0019] Figure 2 is a schematic block diagram illustrating another embodiment of a grid system in accordance with the present invention;

[0020] Figure 3 is a schematic block diagram illustrating one embodiment of a global autonomic manager in accordance with the present invention;

[0021] Figure 4 is a schematic block diagram illustrating one embodiment of a client in accordance with the present invention;

[0022] Figure 5 is a schematic flow chart diagram illustrating one embodiment of a stimulus response method in accordance with the present invention;

[0023] Figure 6 is a schematic flow chart diagram illustrating one embodiment of an autonomic regulation method in accordance with the present invention; and

[0024] Figure 7 is a schematic flow chart diagram illustrating one embodiment of an autonomic prediction method in accordance with the present invention.

DETAILED DESCRIPTION OF THE INVENTION

[0025] Many of the functional units described in this specification have been labeled as modules, in order to more particularly emphasize their implementation independence. For example, a module may be implemented as a hardware circuit comprising custom VLSI circuits or gate arrays, off-the-shelf semiconductors such as logic chips, transistors, or other discrete components. A module may also be implemented in programmable hardware devices such as field programmable gate arrays, programmable array logic, programmable logic devices or the like.

[0026] Modules may also be implemented in software for execution by various types of processors. An identified module of executable code may, for instance, comprise one or more physical or logical blocks of computer instructions which may, for instance, be organized as an object, procedure, or function. Nevertheless, the executables of an identified module need not be physically located together, but may comprise disparate instructions stored in different locations which, when joined logically together, comprise the module and achieve the stated purpose for the module.

[0027] Indeed, a module of executable code could be a single instruction, or many instructions, and may even be distributed over several different code segments, among different programs, and across several memory devices. Similarly, operational data may be identified and illustrated herein within modules, and may be embodied in any suitable form and organized within any suitable type of data structure. The operational data may be collected as a single data set, or may be distributed over different locations including over different storage devices, over disparate memory devices, and may exist, at least partially, merely as electronic signals on a system or network.

[0028] Furthermore, modules may also be implemented as a combination of software and one or more hardware devices. For instance, a module may be embodied in the combination of a software executable code stored on a memory device. In a further example, a module may be the combination of a processor that operates on a set of operational data.

Still further, a module may be implemented in the combination of an electronic signal communicated via transmission circuitry.

[0029] Reference throughout this specification to “one embodiment,” “an embodiment,” or similar language means that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one embodiment of the present invention. Thus, appearances of the phrases “in one embodiment,” “in an embodiment,” and similar language throughout this specification may, but do not necessarily, all refer to the same embodiment.

[0030] Furthermore, the described features, structures, or characteristics of the invention may be combined in any suitable manner in one or more embodiments. In the following description, numerous specific details are provided, such as examples of programming, software modules, user selections, network transactions, database queries, database structures, databases, hardware modules, hardware circuits, hardware chips, etc., to provide a thorough understanding of embodiments of the invention. One skilled in the relevant art will recognize, however, that the invention can be practiced without one or more of the specific details, or with other methods, components, materials, and so forth. In other instances, well-known structures, materials, or operations are not shown or described in detail to avoid obscuring aspects of the invention.

[0031] Figure 1 depicts a grid system 100 that comprises a grid server 102 connected to multiple clients 104-110, or nodes, via a communications channel 112. The illustrated grid system 100 is similar to a local area network (LAN), and the communications channel 112 may be, in one embodiment, an Ethernet communications channel, a wireless communications channel, or another equivalent communications channel. Likewise, the communications channel 112 may comprise a combination of various types of communications channels. Although the depicted grid system 100 includes one grid server 102 and four clients 104-110, the grid system 100 may comprise a combination of various network configurations having fewer or more clients 104-, more than one server 102, or

alternate server configurations. In a further embodiment, the grid system 100 also may include a subscription manager (not shown) as described with reference to Figure 2. In one embodiment, the grid server 102 may concurrently act as the subscription manager of the grid system 100.

[0032] The grid system 100 is configured, in one embodiment, to execute a grid application. A grid application is a collection of work items that together achieve a specified objective. For example, a grid application may determine very complex mathematical calculations, including weather forecasting, stock market development, and so forth. A grid application also may process large-scale multimedia operations. In another embodiment, a grid application may perform data backup operations on large and diverse amounts of data. In each of these scenarios, execution of a grid application may require the cooperation of several nodes 104-110 within the grid system 100.

[0033] A grid application may be divided into jobs, or single units of work. The several jobs of a grid application may be executed concurrently, serially, or co-dependently on one or more of the various nodes 104-110. Each of the nodes 104-110 may allocate certain performance resources to the grid system 100 for execution of grid applications. These performance resources made available by the clients 104-110 may include processor capability, processor capacity, storage capacity, memory capacity, and other similar resources. In one embodiment, a client 104-110 may dedicate a specific amount of total processor capability, storage capacity, or memory capacity to the grid system 100 for execution of grid applications.

[0034] Each client 104-110 may act as either a source client or a resource client, depending on the role of the client 104-110 in a particular grid application. For example, where the client 104-110 initiates a grid application, the client 104-110 acts as a source client. Alternately, where the client 104-110 makes local performance resources available for execution of a remotely initiated grid application, the client 104-110 acts as a resource client. For example, in the case of a grid backup operation, a source client may have backup

data files on one or more resource clients where the resource clients allocate some available storage to the grid system 100 for such backup grid applications. In a further embodiment, the grid server 102 also may act as a client on the grid system 100 in that it may initiate grid applications and make local performance resources available to the grid system 100, in addition to acting as a server 102 for the network 100.

[0035] Figure 2 depicts another embodiment of a grid system 200 that is similar in some aspects to the grid system 100 of Figure 1. The illustrated grid system 200 operates over the internet 202, which provides a communications channel among the various other components of the grid system 200. The illustrated grid system 200 also includes network systems 204, 206, which are similar to the grid system 100 shown in Figure 1, that form sub-systems within the grid system 200 of Figure 2. Additionally, the grid system 200 may include other clients 208, 210 that are directly connected to the internet in that they are not a part of a local network.

[0036] The grid system 200 also may include a subscription manager 212 configured to manage a client subscription to the grid computing system 200. The subscription manager, in one embodiment, may manage the use of the grid system 100 by a subscribed client in terms of client fees or permission for a client to use a grid system resource or expect a certain level of service from the grid computing system 100. The subscription manager 212 may alternatively be connected to other network systems 204, 206 within the grid system 200. In a further embodiment, the grid system 200 may have multiple subscription managers 212 that each manages independently defined subscription groups.

[0037] As mentioned above, other similar grid system configurations may be employed in place of or in addition to the grid systems 100, 200 depicted in Figures 1 and 2. In the following description, reference to either of the grid systems 100, 200 is meant to interchangeably refer to either or both of the grid systems 100, 200, unless the exclusion of one of the grid systems 100, 200 is explicitly noted.

[0038] Figure 3 depicts one embodiment of a global autonomic manager 300. The illustrated global autonomic manager 300 is configured, in one embodiment, to facilitate autonomic control of performance resources on the grid system 100. In one embodiment, the global autonomic manager 300 includes a central processing unit (CPU) 302, a local storage device 304, a user interface 306, a network interface 308, a memory 310, and a global autonomic management apparatus 312. The CPU 302 is configured generally to execute operations within the global autonomic manager 300. The user interface 306, in one embodiment, is configured to allow a user to interact with the global autonomic manager 300, including allowing input data and commands from a user and communicating output data to the user. The network interface 308 is configured, in one embodiment, to facilitate network communications of the global autonomic manager 300 over the communications channel 112 of the grid network 100.

[0039] The local memory 310 is configured, in one embodiment, to store several data and metadata files that may be used in conjunction with the autonomic control of grid performance resources. In an alternative embodiment, some or all of these data and metadata files may be replicated in the local storage device 304. In a further embodiment, one or all of these data and metadata files may be stored exclusively in the local storage device 304 rather than in the memory 310. Similarly, these data and metadata files may be stored on a combination of local memory 310 and storage 304. In another embodiment, one or all of these data and metadata files may be stored in distributed storage on the grid system 100. Although the present description refers to “files,” the present invention is understood to operate in substantially the same manner using other electronic memory and storage structures. Reference herein to a data file or metadata file is understood to equivalently refer to other such electronic memory and storage structures.

[0040] In particular, the memory 310 may store a system resource profile 314, one or more system prediction policies 316, one or more system regulation policies 318, one or more system termination policies 320, and one or more system arbitration policies 321. The

system resource profile 314, in one embodiment, is configured to store one or more parameters indicating the specific performance resources dedicated or allocated to the grid system 100. The system prediction policies 316, in one embodiment, are configured to define one or more prediction trigger events and indicate what action the global autonomic manager 300 will take in response to a recognized prediction trigger event.

[0041] Likewise, the system regulation policies 318, in one embodiment, are configured to define one or more regulation trigger events and indicate what action the global autonomic manager 300 will take in response to a recognized regulation trigger event. Similarly, the system termination policies 320, in one embodiment, are configured to define one or more termination trigger events and indicate what action the global autonomic manager 300 will take in response to a recognized termination trigger event. Furthermore, the system arbitration policies 321 are configured, in one embodiment, to store one or more arbitration policies that indicate what action the global autonomic manager 300 will take in response to a recognized conflict.

[0042] A trigger event may include any action that is defined by one of the policies 316, 318, 320, 321 described above. For example, a trigger event may comprise a change in data flow, a change in time, a scheduled time, a grid operation, a client operation, an automatic operation, a manual operation, a user request, or any of a variety or combination of these and other types of client, grid, user, and other monitored actions.

[0043] The global autonomic management apparatus 312 is configured, in one embodiment, to facilitate autonomic control of performance resources on the grid system 100. The illustrated global autonomic management apparatus 312 includes a global monitor module 322, a global policy module 324, and a global regulation module 326.

[0044] In one embodiment, the global monitor module 322 is configured to monitor the grid system 100 for a trigger event defined by one of the system prediction, regulation, or termination policies 316, 318, 320. The global policy module 324, in one embodiment, is

configured to allow a user to access and define the system prediction, regulation, or termination policies 316, 318, 320.

[0045] The global regulation module 312 is configured, in one embodiment, to autonomically regulate the performance resources on the grid system 100 in response to a recognized trigger event and according to the system prediction, regulation, or termination policies 316, 318, 320. The depicted global regulation module 326 includes a global reservation module 328, a global termination module 330, a global arbitration module 332, and a global profile module 334.

[0046] The global reservation module 328, in one embodiment, is configured to reserve a performance resource on the grid system 100 for a specific grid application or grid system operation. In a similar manner, the global termination module 330 is configured, in one embodiment, to terminate the reservation of a performance resource for a specific grid application. In this way, the global autonomic management apparatus 312 may assign and unassign certain allocated performance resource on the grid system 100 for execution of specific grid applications or jobs.

[0047] In one embodiment, the global arbitration module 332 is configured to arbitrate among competing grid application requirement and conflicting performance resource assignments. For example, the global arbitration module 332 may temporarily suspend assignment or execution of a job that is assigned to a performance resource that is currently in use. In a further embodiment, the global arbitration module may be configured to arbitrate between conflicting grid system operations and local client operations. It is possible, for example, that a client 104-110 may reclaim a performance resource that was previously assigned to be used by the grid system 100. In this case, it may be necessary to re-execute the aborted grid operation or to assign alternative performance resources to complete the suspended grid operation.

[0048] The global profile module 334, in one embodiment, may be configured to manage the system resource profile 314. In one embodiment, the global profile module 334

may track the available performance resources throughout the grid system 100. In a further embodiment, the global profile module 334 also may track the usage of available performance resources on the grid system 100. Additionally, the global profile module 334 may be configured to synchronize the system resource profile 314 with a similar profile on one or more of the clients 104-110.

[0049] Figure 4 depicts one embodiment of a client 400 that may operate as either a source client or a resource client within the grid system 100. Like the global autonomic manager 300 of Figure 3, the client 400 includes a CPU 402, a local storage device 404, a user interface 406, a network interface 408, and a memory 410. The illustrated client 400 also includes a local autonomic management apparatus 412. The CPU 402, user interface 406, and network interface 408 of the client 400 are substantially similar to the CPU 302, user interface 306, and network interface 308 of the global autonomic manager 300.

[0050] The memory 410, in one embodiment, may be configured to store a client resource profile 414, one or more client prediction policies 416, one or more client initiation policies 418, one or more client regulation policies 420, and one or more client termination policies 422. The client resource profile 414 may be similar to the system resource profile 314 of the global autonomic manager 300. However, the client resource profile 414 is configured to store one or more parameters indicating the specific performance resources dedicated to the grid system 100 by the client 400. For example, the client may dedicate bandwidth, storage, memory, processor cycles, input/output devices, and so forth to be used for a grid application.

[0051] The client prediction policies 416, client regulation policies 420, and client termination policies 422 may be similar to the system prediction policies 316, system regulation policies 318, and system termination policies 320, respectively, on the global autonomic manager 300. For example, the client prediction policies 416 may include a policy to request that the global autonomic management apparatus 312 reserve grid

performance resources in anticipation of a grid application originating from the requesting client 400.

[0052] However, these client policies 416, 420, 422 may be specific to the client 400 in terms of the client performance resources dedicated to the grid system 100. For example, the client regulation policies 420 may include policies that regulate the allocation of client performance resources according to the amount of grid operations originating from the client 400. In a certain embodiment, if the client 400 is initiating a substantial amount of grid system operations, the client 400 may allocate additional performance resources to the grid system 100 according to a client regulation policy 420.

[0053] The client initiation policies 416, in one embodiment, include policies that indicate what action the client 400 will take in response to a given initiation trigger event. For example, the client initiation policies 416 may include policies that regulate when a grid application is initiated. For instance, a client initiation policy 316 may specify that grid applications are only initiated after working hours, during lunch hour, when the client processor is idle for a specified time period, and so forth.

[0054] The illustrated local autonomic management apparatus 412, in one embodiment, may be configured to facilitate autonomic control of performance resources on the grid system 100. The depicted local autonomic management apparatus 412 includes a client monitor module 424, a client policy module 426, a client notification module 428, and a client regulation module 430.

[0055] In one embodiment, the client monitor module 424 is configured to monitor the client 400 for a trigger event defined by one of the client prediction, initiation, regulation, or termination policies 416, 418, 420, 422. The client policy module 426, in one embodiment, is configured to allow a user to access and define the client prediction, initiation, regulation, or termination policies 416, 418, 420, 422. The client notification module 428 is configured, in one embodiment, to notify the global autonomic manager 300 of

pertinent activity on the client 400, such as modification of the client resource profile 414, initiation of a grid application, termination of a grid application, and so forth.

[0056] The client regulation module 430 is configured, in one embodiment, to autonomically regulate the performance resources on the client 400 in response to a recognized trigger event and according to the client prediction, initiation, regulation, or termination policies 416, 418, 420, 422. The depicted client regulation module 430 includes a client allocation module 432, a client reclamation module 434, a client initiation module 436, a client termination module 438, and a client profile module 440.

[0057] The client allocation module 432 is configured, in one embodiment, to allocate a client performance resource to the grid system 100 according to the client resource profile 414. As mentioned above, the client performance resource may be any performance resource that the client 400 may make available for grid system operations, including client accessibility, client bandwidth, client processor, client storage, client memory, and so forth. In a similar manner, the client reclamation module 434, in one embodiment, is configured to reclaim a client performance resource that is allocated to the grid system 100. Reclaiming a client performance resource makes the resource unavailable to the grid system 100. Hence, the global autonomic management apparatus 312 cannot reserve, via the global reservation module 328, a reclaimed client performance resource. In this way, the local autonomic management apparatus 412 may allocate and reclaim performance resources on the grid system 100, for example, when such performance resources are not being utilized by the client 400.

[0058] The client initiation module 436, in one embodiment, is configured to initiate a grid application on the grid system 100 in response to an initiation trigger event and according to the client initiation policies 418. Similarly, the client termination module 438 is configured, in one embodiment, to terminate a grid application initiated by the client 400 in response to a termination trigger event and according to the client termination policies 422.

[0059] The client profile module 440, in one embodiment, may be configured to manage the client resource profile 414. In one embodiment, the client profile module 440 may track the allocated client performance resources. In a further embodiment, the client profile module 440 also may track the usage of allocated client performance resources. Additionally, the client profile module 440 may be configured to synchronize the client resource profile 414 with the system resource profile 314 on the global autonomic manager 300.

[0060] The following schematic flow chart diagrams that follow are generally set forth as logical flow chart diagrams. As such, the depicted order and labeled steps are indicative of one embodiment of the presented process. Other steps and processes may be conceived that are equivalent in function, logic, or effect. Additionally, the format and symbology employed are provided to explain the logical steps of the process and are understood not to limit the scope of the process. Likewise, although various arrow types and line types may be employed in the flow chart diagrams, they are understood not to limit the scope of the corresponding process. Indeed, some arrows or other connectors may be used to indicate only the logical flow of the process. For instance, an arrow may indicate a waiting or monitoring period of unspecified duration between enumerated steps of the depicted process.

[0061] Figure 5 depicts one embodiment of an autonomic initiation method 500 that may be employed by the local autonomic management apparatus 412 on the client 400. The illustrated autonomic initiation method 500 begins 502 as the client 400 monitors 504 for an initiation trigger event. In one embodiment, the client 400 may employ the client monitor module 424 to monitor for an initiation trigger event. The initiation trigger event, in one embodiment, is defined by one of the client initiation policies 418.

[0062] The autonomic initiation method 500 continues 506 to monitor 504 for an initiation trigger event until the client 400 recognizes 508 a defined initiation trigger event. The client 400 then initiates the appropriate grid application as indicated by the client

initiation policies 418. In one embodiment, the client 400 employs the client initiation module 436 to initiate the grid application and notifies the global autonomic manager 300 via the client notification module 428. The client 400 then executes the grid application over the grid system 100 until the grid application is complete 514. The depicted autonomic initiation method 500 then ends 516.

[0063] Figure 6 depicts one embodiment of an autonomic regulation method 600 that may be employed by either the global autonomic management apparatus 312 on the global autonomic manager 300 or the local autonomic management apparatus 412 on the client. The illustrated autonomic regulation method 600 begins 602 in the same was as the autonomic initiation method 500 described above. In fact, the autonomic regulation method 600, in one embodiment, follows substantially the same steps 604, 606, 608 to initiate 610 the grid application as the autonomic initiation method 500.

[0064] However, while the grid application operations are being executed 612 on the grid system 100 and before the grid application is completed 614, the autonomic regulation method 600 monitors the grid system 100 for a regulation trigger event. In one embodiment, the client 400 may employ the client monitor module 424 to monitor for a regulation trigger event. In a further embodiment, the global autonomic manager 300 may employ the global monitor module 322 to monitor for a regulation trigger event. The regulation trigger event, in one embodiment, is defined by one of the client regulation policies 420. Alternately, the regulation trigger event may be defined by one of the system regulation policies 318.

[0065] If a regulation trigger event is not recognized 618, the grid application continues to execute normally. However, if a regulation trigger event is recognized 618, the autonomic regulation method 600 regulates 620 the performance resources according to the corresponding client regulation policies 420 or system regulation policies 318. The autonomic regulation method 600 may regulate 620 the performance resources, in one embodiment, by reserving additional performance resources for the grid application or by terminating existing performance resource reservations. Alternately, the autonomic

regulation method 600 may regulate 620 the performance resources by allocating or reclaiming client performance resources, or even by terminating the grid application via the client termination module 438.

[0066] The grid system 100 continues to execute the grid application, subject to the autonomic regulations invoked, until the grid application is either terminated by the client 400 or global autonomic manager 300, or until the grid application is complete 614. The depicted autonomic regulation method 600 then ends 622.

[0067] Figure 7 depicts one embodiment of an autonomic prediction method 700 that may be employed by either the global autonomic management apparatus 312 on the global autonomic manager 300 or the local autonomic management apparatus 412 on the client. The autonomic prediction method 700 is invoked, in one embodiment, prior to the initiation of a grid application. In a further embodiment, the autonomic prediction method 700 may be invoked prior to recognizing 508, 608 an initiation trigger event, as described with reference to Figures 5 and 6.

[0068] The illustrated autonomic prediction method 700 begins 702 by monitoring 704 for a prediction trigger event. In one embodiment, the client 400 may employ the client monitor module 424 to monitor for a prediction trigger event. In a further embodiment, the global autonomic manager 300 may employ the global monitor module 322 to monitor for a prediction trigger event. The prediction trigger event, in one embodiment, is defined by one of the client prediction policies 416. Alternately, the prediction trigger event may be defined by one of the system prediction policies 316.

[0069] The autonomic prediction method 700 continues 706 to monitor 704 for a prediction trigger event until the client 400 recognizes 708 a defined prediction trigger event. The client 400 then regulates 710 the performance resources of the client 400 or grid system 100 as indicated by the client prediction policies 416 or the system prediction policies 316, respectively. The autonomic prediction method 700 may regulate 710 the performance resources, in one embodiment, by reserving additional performance resources for the grid

application or by terminating existing performance resource reservations. Alternately, the autonomic prediction method 700 may regulate 710 the performance resources by allocating or reclaiming client performance resources via the client allocation module 432 and client reclamation module 434, respectively. Additionally, the client 400 may notify the global autonomic manager 300 of changes to performance resource allocations and the client resource profile 414 via the client notification module 428. The depicted autonomic prediction method 700 then ends 712.

[0070] As one example of the foregoing methods 500, 600, 700, the global autonomic manager 300 may monitor the performance of the grid applications on the grid computing system 100. While monitoring the performance of grid, if the overall performance begins to slow down, the global regulation module 326 may invoke system regulation policy 318 to increase the amount of performance resources allocated to the grid application.

[0071] For example, a movie studio may perform a rendering project that is due on a particular date. To meet this deadline, the project managers or business executives may decide to increase the client performance resources available to the grid system 100. In one embodiment, increasing the performance resources for one grid application may reduce the performance resources available to other clients and grid applications on the grid system 100. In this way, the regulation policy 318 defines the action to be taken in response to the request for increased resources. Additionally, the global autonomic management apparatus 312 helps the global autonomic manager 300 and the client 400 to manage the grid application in order to meet the project deadline. In an alternate embodiment, the request for increased performance resources may occur autonomically according to regulation policies 318 defining project constraints and deadlines.

[0072] In another embodiment, a high level regulation policy 318 may be defined to alter the performance resources available to the grid system 100 and to the client 400 based on the overall performance of the grid system 100. For instance, using the same rendering

example discussed above, a regulation policy 318 may allow client performance resources to be dynamically reclaimed by the client 400 if the rendering grid application is not currently making use of the allocated client performance resources. Similarly, the reclaimed client performance resources may be reallocated to the rendering grid application when they are requested to perform additional rendering grid application operations. This is one embodiment of dynamic regulation policy 318.

[0073] Many other autonomic prediction, initiation, regulation, termination, and arbitration policies may be employed in a variety of grid systems 100 and application environments. For instance, the policy may define an on/off control. That is, if the demand for its service is required it is turned on, otherwise it is off. In another embodiment, the policy may define a proportional control. That is, if the demand has increased or decreased, the resources may be allocated or reclaimed in proportion to the demand. In another embodiment, the policy may be time delayed or phased, according to a step change in demand. For example, the performance resources may be gradually allocated or reclaimed over some period of time until the proper level or resource availability and allocation is achieved.

[0074] With regard to the prediction policies, the global autonomic manager 300 may collect and use historical or other information to predict when a particular load, resource availability, resource allocation, etc., may need to change. For example, the local processing load on a client 400 may adversely affect performance resource availability on the grid system 100 at a certain time each day. A predictive policy allows the global autonomic manager 300 and client 400 to predictively adjust grid application resource availability in anticipation of a typical resource usage. This might include delaying the grid application or causing the grid application to execute sooner than normal. The particular prediction, initiation, regulation, termination, and arbitration policies are not limited to any particular policy, control algorithm, etc.

[0075] In one embodiment, the change in policies at the global autonomic manager 300 or at the client 400 may be coupled to a subscription manager 212. The subscription manager 212, in one embodiment, may adjust the fees depending on the change in service. In another embodiment, that subscription manager 212 may prevent specific or global changes in policies. For example, a business may subscribe to a grid system 100 and choose to contribute some level of performance resources. The fee charged to the business may be proportional to the level of grid application demand and level of performance resources contributed. However, certain upper and/or lower thresholds may be fixed, corresponding with particular financial or other business criteria. In a further embodiment, the subscription manager 212 may block potential changes in one or more policies according to these thresholds.

[0076] With further regard to the subscription manager 212, the subscription manager 212, in one embodiment, is an apparatus for managing the information collected, used, or generated in the process of determining user fees, controlling the level of service, controlling the use of the service, controlling the contribution of performance resources, etc. to or for a grid application, from or to a customer, business, etc.

[0077] In one embodiment, the subscription manager 212 may serve at least two purposes. First, it may determine the user fees to be charged to a user based on usage of the grid resources by the user and/or contribution of performance resources by the user to the grid. Second, the subscription manager 212 may control the access, use, level of use, and so forth, to the grid system 100 and grid resources. The subscription manager 212 also may control the allocation, level of contribution, and so forth, of client performance resources to the grid system 100 based on autonomic policies described herein.

[0078] In order to manage the subscriptions of various clients 400 to the grid system 100, the subscription manager 212 may create and store a client profile, a global profile, and a customer profile. In one embodiment, the global profile of the subscription manager 212 may contain information regarding performance resource allocation and usage in order to

determine the user fee for a specific customer. In one embodiment, the global profile of the subscription manager 212 may be generic to all performance resources and clients 400 using the grid system 100.

[0079] In one embodiment, the customer profile contains information that relates the global profile to the particular customer. The customer profile may aggregate information about a particular customer, including information about client performance resource allocation and locally invoked grid applications. The customer profile may be used to determine the overall fee that a customer is charged. Similarly, in one embodiment, the client profile in the subscription manager 212 may contain similar information that corresponds to a specific client 400.

[0080] In one embodiment, the subscription manager 212 determines user fees based on one or more of the instantaneous, average, maximum, minimum, planned, reserved, peak, and so forth, use of the grid system 100 by client 400 for a grid application. In another embodiment, the subscription manager 212 may track the allocation of client performance resources to the grid system 100 by a client 400. The subscription manager 212 may track one or more of the instantaneous, average, maximum, minimum, planned, reserved, peak, and so forth, level contributed. In a further embodiment, the subscription manager 212 track a combination of one or more of the factors listed above.

[0081] In another embodiment, the subscription manager 212 may monitor and control the execution of an autonomic policy by a global autonomic manager 300 or the client 400. For example, a business may subscribe to a grid system 100 for a backup retrieve grid application. To keep costs down, the business may decide to contribute performance resources to the grid system 100 from each of the connected clients 400. If a user decides to reclaim the allocated performance resources of a particular client and reduce his contribution to zero, the subscription manager 212 may alter the client profile and customer profile to determine the appropriate fee. According to the global profile of the subscription manager 212, the global autonomic manager 300 of the grid system 100 may maintain upper and

lower thresholds for performance resource allocation, thereby preventing such a reclamation of all allocated resources.

[0082] In another embodiment, the subscription manager 212 may control a policy change requested by a client 400 or by a global autonomic manager 300. The customer profile of the subscription manager 212 may prevent certain changes to the resource allocation or to the grid application usage of the client 400. For example, the client profile may have a limit on the total cost that a customer may incur in a predetermined billing period. The subscription manager 212 may block certain uses by a client 400 if these limits are exceeded.

[0083] The present invention may be embodied in other specific forms without departing from its spirit or essential characteristics. The described embodiments are to be considered in all respects only as illustrative and not restrictive. The scope of the invention is, therefore, indicated by the appended claims rather than by the foregoing description. All changes which come within the meaning and range of equivalency of the claims are to be embraced within their scope.

[0084] What is claimed is: